



Please type a plus sign (+) inside this box → +

HDP/SB/21 based on PTO/SB/21 (08-00)

AF&  
1/2/05

# TRANSMITTAL FORM

(to be used for all correspondence after initial filing)

Application Number	09/620,053
Filing Date	July 20, 2000
Inventor(s)	Yang CAO
Group Art Unit	2661
Examiner Name	Ian N. Moore
Attorney Docket Number	29250-000971/US

## ENCLOSURES (check all that apply)

<input checked="" type="checkbox"/> Fee Transmittal Form  <input checked="" type="checkbox"/> Fee Attached  <input type="checkbox"/> Amendment  <input type="checkbox"/> After Final  <input type="checkbox"/> Affidavits/declaration(s)  <input type="checkbox"/> Extension of Time Request  <input type="checkbox"/> Express Abandonment Request  <input type="checkbox"/> Information Disclosure Statement  <input type="checkbox"/> Certified Copy of Priority Document(s)  <input type="checkbox"/> Response to Missing Parts/ Incomplete Application  <input type="checkbox"/> Response to Missing Parts under 37 CFR 1.52 or 1.53	<input type="checkbox"/> Assignment Papers (for an Application)  <input type="checkbox"/> Letter to the Official Draftsperson and _____ Sheets of Formal Drawing(s)  <input type="checkbox"/> Licensing-related Papers  <input type="checkbox"/> Petition  <input type="checkbox"/> Petition to Convert to a Provisional Application  <input type="checkbox"/> Power of Attorney, Revocation Change of Correspondence Address  <input type="checkbox"/> Terminal Disclaimer  <input type="checkbox"/> Request for Refund  <input type="checkbox"/> CD, Number of CD(s) _____	<input type="checkbox"/> After Allowance Communication to Group  <input type="checkbox"/> LETTER SUBMITTING APPEAL BRIEF AND APPEAL BRIEF (w/clean version of pending claims)  <input checked="" type="checkbox"/> Appeal Communication to Group (Notice of Appeal, <u>Brief</u> , Reply Brief)  <input type="checkbox"/> Proprietary Information  <input type="checkbox"/> Status Letter  <input type="checkbox"/> Other Enclosure(s) (please identify below):		
<table><tr><td>Remarks</td><td></td></tr></table>			Remarks	
Remarks				

## SIGNATURE OF APPLICANT, ATTORNEY, OR AGENT

Firm or Individual name	Harness, Dickey & Pierce, P.L.C.	Attorney Name	John E. Curtin	Reg. No.	37,602
Signature					
Date	September 16, 2005				

--



Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PTO/SB/17 (12-04)  
Approved for use through 07/31/2006. OMB 0651-0032  
U.S. Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

# FEE TRANSMITTAL for FY 2005

Effective 10/01/2004. Patent fees are subject to annual revision.

☐ Applicant claims small entity status. See 37 CFR 1.27

TOTAL AMOUNT OF PAYMENT (\$ ) 500

## Complete if Known

Application Number	09/620,053
Filing Date	July 20, 2000
First Named Inventor	Yang CAO
Examiner Name	Ian N. Moore
Art Unit	2661
Attorney Docket No.	29250-000971/US

## METHOD OF PAYMENT (check all that apply)

☒ Check ☐ Credit card ☐ Money ☐ Other ☐ None  
Order

☒ Deposit Account:

Deposit  
Account  
Number

08-0750

Deposit  
Account  
Name

Harness, Dickey & Pierce, PLC

The Director is authorized to: (check all that apply)

☐ Charge fee(s) indicated below ☒ Credit any overpayments  
☒ Charge any additional fee(s) during the pendency of this application  
☐ Charge fee(s) indicated below, except for the filing fee  
to the above-identified deposit account.

## FEE CALCULATION

### 1. BASIC FILING FEE

Large Entity		Small Entity		Fee Description	Fee Paid
Fee Code	Fee (\$)	Fee Code	Fee (\$)		
1011	300	2011	150	Utility filing fee	
1012	200	2012	100	Design filing fee	
1013	200	2013	100	Plant filing fee	
1014	300	2014	150	Reissue filing fee	
1005	200	2005	100	Provisional filing fee	

SUBTOTAL (1)

(\$ ) 0

### 2. EXTRA CLAIM FEES FOR UTILITY AND REISSUE

		Extra Claims	Fee from below	Fee Paid
Total Claims	-20 **	= 0	X	= 0
Independent Claims	-3 **	= 0	X	= 0
Multiple Dependent				= 0

Large Entity		Small Entity		Fee Description
Fee Code	Fee (\$)	Fee Code	Fee (\$)	
1202	50	2202	25	Claims in excess of 20
1201	200	2201	100	Independent claims in excess of 3
1203	360	2203	180	Multiple dependent claim, if not paid
1204	200	2204	100	** Reissue independent claims over original patent
1205	50	2205	25	** Reissue claims in excess of 20 and over original patent

SUBTOTAL (2)

(\$ ) 0

\*\*or number previously paid, if greater; For Reissues, see above

## FEE CALCULATION (continued)

### 3. ADDITIONAL FEES

Large Entity		Small Entity		Fee Description	Fee Paid
Fee Code	Fee (\$)	Fee Code	Fee (\$)		
1051	130	2051	65	Surcharge - late filing fee or oath	
1052	50	2052	25	Surcharge - late provisional filing fee or cover sheet.	
1053	130	1053	130	Non-English specification	
1812	2,520	1812	2,520	For filing a request for reexamination	
1804	920*	1804	920*	Requesting publication of SIR prior to Examiner action	
1805	1,840*	1805	1,840*	Requesting publication of SIR after Examiner action	
1251	120	2251	60	Extension for reply within first month	
1252	450	2252	225	Extension for reply within second month	
1253	1020	2253	510	Extension for reply within third month	
1254	1,590	2254	795	Extension for reply within fourth month	
1255	2,160	2255	1080	Extension for reply within fifth month	
1401	500	2401	250	Notice of Appeal	500
1402	500	2402	250	Filing a brief in support of an appeal	
1403	1000	2403	500	Request for oral hearing	
1452	500	2452	250	Petition to revive - unavoidable	
1453	1500	2453	750	Petition to revive - unintentional	
1501	1400	2501	700	Utility issue fee (or reissue)	
1502	800	2502	400	Design issue fee	
1460	130	1460	130	Petitions to the Commissioner	
1807	50	1807	50	Processing fee under 37 CFR 1.17 (q)	
1806	180	1806	180	Submission of Information Disclosure Stmt	
8021	40	8021	40	Recording each patent assignment per property (times number of properties)	
1809	790	2809	395	Filing a submission after final rejection (37 CFR § 1.129(a))	
1810	790	2810	395	For each additional invention to be examined (37 CFR § 1.129(b))	
1801	790	2801	395	Request for Continued Examination (RCE)	

Other fee (specify) \_\_\_\_\_

\*Reduced by Basic Filing Fee Paid SUBTOTAL (3) (\$ ) 500

### 4. SEARCH/EXAMINATION FEES

1111	500	2111	250	Utility Search Fee	
1112	100	2112	50	Design Search Fee	
1113	300	2113	150	Plant Search Fee	
1114	500	2114	250	Reissue Search Fee	
1311	200	2311	100	Utility Examination Fee	
1312	130	2312	65	Design Examination Fee	
1313	160	2313	80	Plant Examination Fee	
1314	600	2314	300	Reissue Examination Fee	

SUBTOTAL (4) (\$ ) 0

## SUBMITTED BY

Name (Print/Type)	John E. Curran	Registration No. (Attorney/Agent)	37,602	Telephone	(703) 668-8000
Signature				Date	September 16, 2005

WARNING: Information on this form may become public. Credit card information should not be included on this form. Provide credit card information and authorization on PTO-2038.



**IN THE U.S. PATENT AND TRADEMARK OFFICE**

Appellant: Yang CAO

Application No.: 09/620,053

Art Unit: 2661

Filed: July 20, 2000

Examiner: Ian N. Moore

For: APPARATUS AND METHOD FOR SYNCHRONOUS AND  
ASYNCHRONOUS SWITCHING OF INTERNET  
PROTOCOL TRAFFIC

Attorney Docket No.: 29250-00971/US

---

**APPELLANT'S BRIEF ON APPEAL UNDER 37 C.F.R. §41.37**

**MAIL STOP APPEAL BRIEF - PATENTS**

Customer Service Window  
Randolph Building  
401 Dulany Street  
Alexandria, VA 22314

September 16, 2005



APPELLANT'S BRIEF ON APPEAL UNDER 37 C.F.R. §41.37

S. Application No.: 09/620,052

Atty. Docket: 29250-000971/US

**TABLE OF CONTENTS**

	<u>Page</u>
APPELLANT'S BRIEF ON APPEAL UNDER 37 C.F.R. §41.37 .....	1
I. REAL PARTY IN INTEREST .....	1
II. RELATED APPEALS AND INTERFERENCES.....	1
III. STATUS OF CLAIMS .....	1
IV. STATUS OF AMENDMENTS .....	1
V. SUMMARY OF CLAIMED SUBJECT MATTER.....	1
i. Overview of the Subject Matter of the Independent Claims .....	1
ii. Additional Text from the Specification in Support of the Claims.....	3
VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL.....	20
VII. ARGUMENTS.....	20
The Section 102 and 103 Rejections.....	20
VIII. CONCLUSION.....	23
APPENDIX A - Claims Appendix	
APPENDIX B - Figure 1	
APPENDIX C - Figure 2	
APPENDIX D - Figure 3	
APPENDIX E - Figure 4	
APPENDIX F - Figure 5	
APPENDIX G - Figure 6	



**APPELLANT'S BRIEF ON APPEAL UNDER 37 C.F.R. §41.37**

**REAL PARTY IN INTEREST:**

The real party in interest in this appeal is Lucent Technologies Inc. Assignment of the application was submitted to the U.S. Patent and Trademark Office on July 20, 2000, and recorded on the same date at Reel 011018, Frame 0738.

**II. RELATED APPEALS AND INTERFERENCES:**

There are no known appeals or interferences that will affect, be directly affected by, or have a bearing on the Board's decision in this Appeal.

**III. STATUS OF CLAIMS:**

Claims 1-42 are pending in the application, with claims 1, 12, 28 and 33 being written in independent form.

Claims 11, 24, 32 and 39 are objected to as being dependent upon a rejected base claim. Claims 1 and 12 remain finally rejected under 35 U.S.C. §102(e), claims 1-10, 12-23, 25-31, 33-38 and 40-42 remain finally rejected under 35 U.S.C. §103(a). Claims 1-42 are being appealed.

**IV. STATUS OF AMENDMENTS:**

A Request for Reconsideration ("Request") was filed on June 24, 2005. In an Advisory Action dated July 8, 2005, the Examiner stated that the Request was considered but did not place the application in condition for allowance.

**V. SUMMARY OF CLAIMED SUBJECT MATTER:**

**(i.) Overview of the Subject Matter of the Independent Claims**

The invention relates to telecommunications systems and, more particularly, to the efficient switching of Internet Protocol (IP) traffic using circuit

and packet switched fabrics, the packet switched fabrics including IP router switch fabrics and asynchronous transfer mode (ATM) switch fabrics.

In accordance with the principles of the present invention a hybrid switch includes packet and circuit switching switch fabrics, a controller, which may be referred to hereinafter as a hybrid switch resource manager, and one or more input/output ports (I/O port). Telecommunications traffic enters the switch and, after the traffic is switched, departs to telecommunications network through the I/O port(s). A new connection admission control (CAC) process is employed by the controller to route Internet Protocol (IP) traffic to either a circuit switch fabric, such as an STM switch fabric, or to packet switch fabric, such as an IP switch fabric. Because the traffic load is shared, in parallel fashion, between the circuit and packet fabrics, neither switch need be of sufficient magnitude to accommodate the entire traffic load (specification, pages 4-5).

In an illustrative embodiment the hybrid switch's controller maps IP traffic into an appropriate ATM service category. IP traffic that is mapped into an ATM service category, such as CBR or rt-VBR, is aggregated "on-the-fly" and switched through the hybrid switch's circuit switched switch fabric. IP traffic that is mapped into another ATM service category is directed to a packet switch fabric for switching in that fabric. As a result, a hybrid switch in accordance with the principles of the present invention provides efficient use of circuit and packet switch fabrics, the overall bandwidth of the switch, and buffers used for access to the switch.

In accordance with the principles of the present invention all IP traffic that is mapped into an ATM CBR service category may be aggregated and switched through a circuit switch fabric. IP traffic that is mapped into ATM rt-VBR service category may be switched through a packet switch fabric in a conventional manner or it may be switched in the same fashion as ATM CBR

traffic. IP traffic mapped into other categories may be switched through the hybrid switch's packet switch fabric (specification, page 5).

**(ii.) Additional Text from the Specification in Support of the Claims**

In eighteen seventy-six, inside a third floor walk-up garret apartment in the Scollay Square section of Boston Massachusetts, only a short distance from the sight of the first battle of the revolutionary war, Alexander Graham Bell spoke the first words transmitted over telephone wires. Bell's transmission of sound over telephone wires initiated a revolution in communications whose scope rivals that of the political revolution initiated by the sound, heard nearby, of "the shot heard round the world" (specification, page 2).

Technical innovations have dramatically transformed the telecommunications industry in the ensuing years. For example, telecommunications switching systems have evolved considerably from "hand operated" systems in which one instrument was electrically connected (through a hierarchical switching network) to another with the intervention of a human operator who would physically plug one circuit into another. Such direct electrical connection of two or more channels between two points (at least one channel in each direction), a connection that provides a user with exclusive use of the channels to exchange information, is referred to as circuit switching, or line switching. Human operators have largely been replaced by systems which employ electronic switching systems (ESS, e.g., 5ESS), in which the instruments are automatically connected through the network by electronic systems.

Nevertheless, such switching systems often still employ circuit switching, a technique which yields highly reliable service, particularly for such "real time" communications applications as voice, in which the momentary loss of a channel is annoying, and repeated such losses are unacceptable (specification, pages 2-3).

Not only has switching technology undergone major changes, the type of traffic being carried on telephone lines has also changed dramatically. Although originally designed for voice traffic and "tuned" to operation in the voice band

between approximately 350 and 4000 Hz, the telecommunications infrastructure also carries data, through the use of various channels, or tones. However, with the growing, use of the Internet, and the potential development such high bandwidth applications such as interactive distance-learning and video on demand,' the existing telecommunications infrastructure is in danger of being overwhelmed. A large portion of the system's transmission medium has been replaced with high speed trunks which employ fiber optic transmission media, microwave media, and line of sight optical media, for example, to meet the ever mounting demand for high speed data transmission capability. Data traffic is increasing at a rate of approximately 300% per year, while voice traffic is only increasing at the relatively slow rate of approximately 5% per year. However, a huge installed base of transmission media, switching devices, and other telecommunications infrastructure provide the telecommunications path for the vast majority of telecommunications providers and users.

Attempts have been made to address quality of service issues that arise, for example, when time-sensitive traffic, such as voice traffic, is transmitted over an internet protocol (IP) network. Since IP was directed, after all, primarily toward the transmission of data, "best effort" transmission was judged to provide sufficient quality for IP transmission and, generally, quality of service was not addressed. However, with the confluence of real-time and non-real-time traffic on networks that may include both circuit and packet switching facilities, conventional approaches to accommodating traffic having varied quality of service requirements under-utilize resources or fail to meet quality of service requirements. Although quality of service categories are available for asynchronous transfer mode (ATM) traffic, the various quality of service categories supported by ATM, and the concomitant varied requirements, exacerbate the difficulty of modeling an ATM compatible switching system. Consequently, many conventional ATM systems either employ conservative systems modeling, which results in underutilized facilities such as bandwidth and/or buffers, or aggressive modeling, which degrade

system performance. A system and method that enable the efficient combination and management of circuit-switched and packet-switched facilities, thereby taking advantage of the tremendous installed base of equipment and facilities while, at the same time, permitting an extensive upgrade of data facilities, which typically employ packet switching systems, would prove advantageous. In particular, a system and method that provides such advantages for IP traffic would be highly desirable (specification, pages 3-4).

The new hybrid switches provided by the present invention may be connected in any of a variety of network topologies. For ease and clarity of description, the illustrative embodiment of Figure 1 (Appendix B) includes six hybrid switches in accordance with the principles of the invention (Node A, Node B, Node C, Node D, Node E and Node F) that are connected in a ring 100, in which Nodes A and B are connected through a link 102, Nodes B and C are connected through a link 104, Nodes C and D are connected through a link 106, Nodes D and E are connected through a link 108, Nodes E and F are connected through a link 110, and Nodes F and A are connected through a link 112. As described in greater detail in the discussion related to Figures 2 (Appendix C) and 3 (Appendix D), each of the hybrid switches (any of the nodes Node A through Node F) includes circuit switch facilities, such as STM switch fabric facilities, packet switch facilities, such as IP switch fabric facilities, Input/Output facilities and a switch manager that, in accordance with principles of the invention, examines incoming traffic and route the traffic to an appropriate switching facility within the hybrid switch. As previously noted, the invention may apply to other network topologies, such as meshes (wherein the nodes A through F may be connected through links 114, 116, 118, 120, 122, and 124, for example). Additionally, within the context of a ring topology, each of the nodes may be connected to a plurality of rings. The illustrative embodiment of Figure 1 was chosen for clarity and simplicity of explanation. Furthermore, although the invention will be discussed in the context of employing a SONET/SDH physical link media, the invention may be used in

conjunction with networks that employ other physical media (specification, pages 6-7).

Each of the Nodes A through F of Figure 1 includes components such as illustrated in the conceptual block diagram of Figure 2. The node 200 includes a controller 202, also referred to herein as a hybrid resource manager 202, a central circuit switch fabric 204, a central packet switch fabric 232, and a plurality of STM units 206, 208, 210, and 212. The hybrid resource manager 202 operates to manage the switch resources through respective communications links 214, 216, 218, 220, and 222 to shelf controllers 224, 226, 228, 230, and a central packet switch fabric 232. The central packet switch fabric in this illustrative embodiment is an asynchronous transfer mode (ATM) switch fabric. The STM units 206, 208, 210 and 212 respectively include input/output (I/O) interfaces 234, 236, 238, and 240, local switch fabrics, 242, 244, 246, and 248, and switch interfaces, 250, 252, 254, and 256. An I/O interface, such as I/O interface 234 may provide a connection, to another network element, or node, through a link such as the link 102 which connects nodes A and B of Figure 1. Traffic arriving at one of the I/O interfaces is routed under control of the hybrid resource manager 202 by a shelf controller, which more directly controls the operation of a local STM switch. Each of the local STM switches may be, for example, an STS-1 level cross-connect. The capacity of the cross-connect may be related to the I/O capacity of the corresponding I/O interface. That traffic which is routed to the central packet switch fabric 232 may be directed through a switch interface, such as SWIF 250, embodied as an advanced "UTOPIA" interface which is capable of transferring both ATM cells and packet-based traffic (specification, pages 7-8).

As described in greater detail in the discussion related to Figures 2, 3 (Appendix D), 4 (Appendix E) and 5 (Appendix F), the hybrid resource manager partitions incoming traffic into streams that are to be switched either by an STM switch fabric or an IP switch fabric. The STM traffic may be switched in a local STM switch, such as local STM switch 242, for example, or it may be switched

through a central STM switch fabric 204 (optional). Typically, a single ring SONET/SDH embodiment would require only one STM unit, such as the unit 206, to provide an interface to the ring and to provide the STM switching function. A multi-ring or mesh connection would include a plurality of the STM units, as illustrated. In such as case, the STM switching may be provided by a central STM switch fabric 204 which could direct traffic into and out of any of the I/O interfaces and to the packet switch fabric 232, as needed. Alternatively, traffic entering one STM unit could be packetized and switched through the central packet switch fabric 232 to travel between I/O interface 234 and I/O interface 240, for example. Each of the I/O interfaces, such as I/O 234, acts as a ring, or line, interface. IP traffic is first mapped into an ATM service category, then switched. The resource manager may perform the IP-to-ATM mapping process, for example. After thus mapping IP into ATM traffic categories; the hybrid resource manager 202 communicates with the next node in a path of which it is a part and determines, "on the fly", which switch fabric, such as a local STM 242, central packet 232 or central circuit 204, will be employed to switch traffic associated with the IP traffic. This decision is based upon the ATM service category into which the IP traffic has been mapped.

An overview of the process by which a hybrid switch services IP traffic in accordance with the principles of the present invention is provided by the flow chart of Figure 3. The process begins in step 300 and proceeds from there to step 302, where various types of IP traffic is mapped into ATM service categories. This mapping process may take place "offline", with the mappings downloaded to the hybrid switch's resource manager or the resource manager may perform the mapping "on the fly", in which case, the mappings maybe changed to accommodate various requirements (specification, page 8-9).

The IP-to-ATM mapping allows the hybrid switch to employ it's connection admission control CAC process to provide access to its switching facilities in much the same way that it would provide access to ATM traffic. As described in

greater detail below, the IP traffic is routed to either an STM or an IP switch fabric within the hybrid switch, depending upon which ATM service category the IP traffic has been mapped into. Although there are a variety of IP reservation protocols, such as resource reservation setup protocol (RSVP), or constraint based routing label distribution protocol (CRLDP), they all contain a traffic description that describes the characteristics of the IP packet flow. A hybrid switch in accordance with the principles of the present invention employs these IP packet flow descriptions as an interface to the CAC process. Although the following illustrative example will be employ details of an RSVP protocol for illustrative purposes, as will be obvious to one skilled in the art, the new hybrid switch may be operated in conjunction with a wide variety of IP reservation protocols, including CRLDP with multi-protocol label switching (MPLS).

An IP signaling protocol, such as RSVP operates as the IP counterpart to the ATM in-band signaling protocol employed to reserve resources within each node along a path through which traffic is to follow. That is, such protocols are employed to reserve resources along the projected path of the IP traffic, just as ATM in-band protocols are employed to reserve resources for ATM traffic prior to carrying the traffic. An IP signaling protocol message, such as an IP RSVP message, will include information that the new hybrid switch may employ to provide IP traffic with a desired quality of service. An IP RSVP message, "Adspec", for example, is transmitted by a sending node to inform receiving nodes of the characteristics of the desired end-to-end communications path. This information is used by the receivers to determine the level of reservation required in order to achieve the desired end to end QoS.

Service categories available to IP traffic include, Guaranteed Service (GS), Controlled Load Service (CLS), and Best Effort (BE). Each of these IP traffic categories may be mapped into an ATM service category, such as CBR, rt-VBR, nrt-VBR, UBR, or ABR. Various mapping procedures are possible and discussed for example in, M. Garrett, M. Borden, Interoperation of Controlled-Load Service

and Guaranteed Service with ATM, <ftp://ftp.isi.edu/in-notes/rfc2381.txt>, which is hereby incorporated by reference (specification, pages 9-10).

An Adspec IP reservation message includes a message header, a Default General Parameters fragment and at least one of the following: Guaranteed Service fragment or Controlled-Load Service fragment. The Default General Parameters fragment includes fields which are updated at each RSVP-capable router along the path in order to present end-to-end values to the receivers. These fields include, Minimum Path Latency, Path Bandwidth, Global Break bit, and Path MTU. The Minimum Path Latency (summation of individual link latencies). represents the end-to-end latency in the absence of any queuing delay. In the case of Guaranteed Service, receivers can add this value to the bounded end-to-end queuing delay to obtain the overall bounded end-to-end delay. The Path Bandwidth is the minimum of individual link bandwidths along the path. The Global Break bit: is cleared when the Adspec is created by the sender. Encountering any routers that do not support RSVP will result in this bit being set to one. And, the Path MTU is the minimum of Maximum Transmission Units of individual links along the path.

For Guaranteed Service, the source Tspec contains peak rate, token bucket rate and bucket depth parameters:  $p_s$ ,  $r_s$  and  $b_s$ . The receiver Tspec contains corresponding parameters  $p_r$ ,  $r_r$  and  $b_r$ . The receiver Rspec also has a rate  $R$  which denote the bandwidth reserved by this receiver. For Controlled Load Service, the source Tspec contains peak rate, token bucket rate and bucket depth parameters:  $p_s$ ,  $r_s$  and  $b_s$ . The receiver Tspec contains corresponding parameters  $p_r$ ,  $r_r$  and  $b_r$ . The receiver Rspec also has a rate  $R$  which denote the bandwidth reserved by this receiver (specification, pages 10-11).

If, for example, IP Guaranteed Service is mapped into the ATM rt-VBR service category, the ATM traffic descriptor parameters (PCR, SCR, MBS) may be set as follows:

$$\text{PCR} = p\_r$$

$$\text{SCR} = r\_r$$

$$\text{MBS} = b\_r$$

Or, if IP Guaranteed Service is mapped into the ATM CBR service category, the value of the ATM PCR may be mapped to the Rspecrate  $r\_r$  if the edge device has sufficient buffer. With little or no burst buffering,  $\text{PCR} = \max(r\_r, p\_r)$ .

IP Controlled Load Service may be mapped into the ATM rt-VBR service category, in which case the ATM traffic descriptor parameters (PCR, SCR, MBS) may be set as follows:

$$\text{PCR} = p\_r$$

$$\text{SCR} = r\_r$$

$$\text{MBS} = b\_r$$

Or, if IP Controlled Load Service is mapped into the ATM ABR service category, the minimal packet rate (MPR) parameter of ABR is set as  $r\_r$ . For Best Effort service there is no traffic description, so it is mapped into ATM LIBR service class.

After mapping IP services into ATM service categories, the process proceeds from step 302 to step 304 where IP traffic is received. From step 304 the process proceeds to step 306 where the arriving IP traffic's ATM "equivalent" category is determined. From step 306 the process proceeds to step 308, where the hybrid switch's CAC is employed to establish a traffic path in a manner set forth in greater detail in the discussions related to Figures 4, 5, and 6 (Appendix G). From step 308, the process proceeds to end in step 310 (specification, page 11).

For each inter-connecting link, such as link 102 between nodes A and B, the total bandwidth of the link is partitioned into a plurality of units and these units are allocated to STM or ATM traffic. If the packet switch fabric 232 supports a UTOPIA-11 interface, the basic unit is an STS-1. If the switch fabric 232 supports a UTOPIA-111 interface, the unit is an STS-3C. For example, if the link 102 is an OC192 link, and the switch fabric supports UTOPIA-III, the link is partitioned into

64 OC-3C based units. For each unit, there is one entry in an ingress and egress resource table, as follows:

Ingress Resource Table:

Tributary No.	Free Flag	Destination Address	Available Bandwidth	Status: Drop/Pass

TABLE 1

Egress Resource Table:

Tributary No.	Free Flag	Destination Address	Available Bandwidth	Status: Drop/Pass

TABLE 2

The destination Address is either based on the IP address for this node or the proprietary address. For each ingress link, there is one ingress resource table associated with it. Correspondingly, there is one egress resource table for each egress link. It is assumed that there is at least one entry in each link's resource table (specification, pages 11-12).

At initialization time, all the free flags are initialized to have the value 0, indicating that the link is free. The available Bandwidth entry takes the initial value of the link's total physical transmission bandwidth, the Destination Address is initialized as 0, and Status is initialized as either Add (egress link) or Drop (ingress link).

At run time the various tributaries' resources may be allocated for different traffic categories with the allocations reflected in the ingress and egress resource tables. For example, resources may be allocated to CBR traffic from tributaries, starting with a top tributary number and working the way down, while other

service categories may be allocated resources from the bottom up. Assume, for example that an egress link is an OC192 link. The link may be partitioned into 64 OC3 tributaries, in which case the link's associated egress table would have 64 entries, one for each tributary (also referred to herein as link traffic units). As traffic requests are received at the node (hybrid switch), resources are allocated, as set forth in greater detail in the discussion related to Figure 4, to service the traffic. An source payload envelope (SPE) mapper may be configured to encapsulate a specific number of channels on the egress link. For example, two OC3C channels provisioned as non-path terminated at the next node, and an OC192 that is provisioned as path terminated at the next node. Additionally, the STS point processor of the inbound OC192 link on the next node would be configured to recognize and process these channels, again, as set forth in greater detail in the discussion related to Figure 4. Traffic received at the "next" node will be processed first by the STS pointer processor in the nodes' I/O shelf, the two OC3C channels may then be served directly by a local STS crossconnect, thus avoiding the necessity of further processing through the UTOPIA interface. The OC192 channel may be sent directly to a centralized IP switch for further processing through a UTOPIA interface (specification, pages 12-13).

A number of ATM service categories, with their attendant quality of service requirements, have been defined and are listed in table 3:

Service Category	QoS
CBR	$PLR \leq \epsilon$ $PDV \leq \delta$ $MaxPTD \leq \gamma$
rt-VBR	$PLR \leq \epsilon$ $PDV \leq \delta$ $MaxPTD \leq \gamma$
nrt-VBR	$PLR \leq \epsilon$
UBR	Best Effort
ABR	$PLR \leq \epsilon$

TABLE 3

Where:

PLR = packet loss rate

PDV = packet delay variation

MaxPTD = maximum end-to-end packet transfer delay and  $\epsilon$ ,  $\delta$  and  $\gamma$  are the corresponding system requirements.

In accordance with the principles of the present invention, traffic falling within the CBR and rt-VBR categories may be treated as "real-time" and switched through a hybrid switch's STM switch fabric. Each incoming ATM connection request is passed to the hybrid resource manager 202 which determines how the traffic is to be switched. In accordance with the principles of the present invention, specific types of ATM requests indicate that the associated ATM traffic is to be switched through a STM switch fabric within the hybrid switch. In this illustrative embodiment, all CBR type ATM traffic is routed by the hybrid resource manager 202 to an STM switch fabric (specification, pages 12-14).

An overview of the process by which the resource manager 202 handles ATM connection requests will be provided in a discussion related to the conceptual block diagram of a three-Node telecommunications system of Figure 4. Each of the

nodes, G through I, is a hybrid telecommunications switch in accordance with the principles of the present invention. Assuming that CBR type ATM traffic originates at Node G and has a destination of Node I STM bandwidth that has already been provisioned for ATM traffic within each of the nodes, and other resources, such as buffer space, is dynamically allocated to the CBR request. For example an ATM request originating at Node G with a destination of Node I would be passed along to Node I and with each subsequent node along the path examining its routing table, as well as the availability of buffers and ingress and egress bandwidth.

Each node's hybrid resource manager could perform this operation. Each node in the path determines whether it has the available resources, whether for pass-through or destination operations, depending upon whether its routing table indicates it is the destination node or not. If each node within the path has the available capacity (for example, ingress and egress bandwidth and buffers), the destination node affirms the ATM request and this affirmation is returned, node to node, to the ingress node (Node G in this example). In this way, resources, including STM switch fabrics within each node, are dynamically provisioned for CBR type ATM traffic in the illustrative network.

The flowchart of Figure 5 illustrates in greater detail the process by which the resource manager 202, through connection admission control, determines the manner in which ATM traffic is to be switched, should the incoming service request indicate that the traffic is CBR traffic. The process begins in step 500, where the resource manager 202 determines that a received connection request is a CBR request. In this illustrative embodiment, all CBR type ATM requests are routed to a STM switch fabric for switching. The process then proceeds to step 502, where the resource manager 202 determines whether the ATM connection request is an "original" CBR request. That is, the resource manager 202 determines whether the node receiving the request, the node associated with the resource manager, is the first node to receive the request, or the request is a pass-through request. If the

request is not an original request, the process branches to step 504, which step will be described below (specification, pages 14-15).

If the request is an "original" request, the process proceeds to step 506, where the resource manager determines whether the switch has sufficient ingress bandwidth to service the request. If there is not sufficient bandwidth for the request, the resource manager denies the request in step 508 and the process proceeds to end in step 510. Bandwidth may be provisioned for different ATM service categories in a variety of ways. For example, all link traffic units, such as OC-3C traffic units, as previously discussed, may initially be "free", and their availability would be reflected by a "0" entry in a corresponding ingress resource table. In response to a connection request having a speed service type, the resource manager determines whether the allocated bandwidth for the specific service category is sufficient for the request and, if bandwidth is available, the resource manager incrementally allocates the bandwidth for the service category. After allocating the bandwidth, the resource manager alters the resource table to reflect the allocation by changing the corresponding "Free Flag" from "0" to "1." Consequently, the resource manager may attempt to allocate the necessary ingress bandwidth, "on the fly" as the request is received and only deny the request if the bandwidth cannot be allocated in the manner just described.

If there is sufficient ingress bandwidth, the process proceeds to step 512 where the resource manager employs the node's routing tables, as previously described in the discussion related to Figure 4, to determine whether the node is the final destination for the related CBR traffic. If the node is the final destination, the process proceeds to step 514 where the resource allocation state is updated, that is, the resource availability information in the resource allocation table is modified to reflect changes, and a request acceptance is returned. From step 514 the process proceeds to end in step 510. If, on the other hand, the node is not the final destination for the CBR traffic, the process proceeds from step 512 to step 516. In step 516, based upon information contained within the CBR request, the

next node and corresponding egress link are computed and the egress link's resource table is located. The process then proceeds to step 518, where the resource manager examines the egress link's corresponding egress resource table. If resources have not yet been allocated for CBR traffic to the destination address, the egress resource table for the link will not have an entry for that address. Therefore, if the resource manager determines that there is no entry within the resource table for the destination address, the process proceeds to step 524, which will be described after step 522 (specification, pages 15-16).

On the other hand, if there is an entry for the destination address, the process proceeds to step 520 where it is determined whether the bandwidth allocated for CBR traffic to the destination address is sufficient for the current request. If there is sufficient bandwidth, the process proceeds to step 522 where the request is accepted, the available bandwidth information is modified to reflect the allocation of bandwidth to the traffic whose request has just been accepted, and the process proceeds to end in step 510. If, in step 520, it is determined that there is insufficient available bandwidth, the process proceeds to step 524. In step 524 the resource manager determines, from the resource table, whether there are "free" resources that may be allocated to CBR traffic for the destination address that are sufficient to handle the CBR request. If there are not, the process proceeds to step 526, where the request is denied, and, from there, to end in step 510. If, on the other hand, it is determined in step 524 that there are sufficient "free" resources to handle the CBR request, the process proceeds to step 528. In step 528, the CBR request is marked as a "passthrough" request and sent to the next node, which may or may not be the destination node. The resource manager then awaits a reply from the next node indicating whether the remaining node(s) in the path, including the egress node, accept the CBR request.

If the response from the remaining nodes is a request rejection, or timeout, the process proceeds to step 526 and from there as previously described. If the response is an acceptance from the other nodes, the process proceeds to step

532, where the resource manager accepts the CBR request. Additionally, the resource manager updates the egress resource table entry by setting the Free Flag to "1", sets the destination address entry as the destination address of the CBR request, modifies the available bandwidth entry by subtracting the bandwidth required for the CBR traffic request just accepted and sets the status as "Add". The CBR bandwidth that is subtracted from the table's available bandwidth entry is indicated by the PCR parameter of the CBR service request. From step 532 the process proceeds to end in step 510 (specification, pages 16-17).

Returning to step 504, the resource manager determines whether the resource table indicates that resources are free to service the incoming CBR traffic and, if not, the process proceeds to step 534 where the node rejects the request and, from there, to end in step 510. If the resource entry indicates there are resources available for the requesting traffic, the process proceeds to step 536, where the resource manager determines whether the node is the destination of the CBR connection. If the node is the destination node the process proceeds to step 538, where the tributary, such as an OC3c, is provisioned as a path termination tributary, the Free Flag of the corresponding ingress resource table is set to "1" and the status is updated to "Drop", indicating that this node is the destination, or "drop" node for the corresponding traffic. The resource manager also returns an acceptance of the request and the process proceeds to end in step 510.

If the node is not the destination node for the CBR request, the process proceeds from step 536 to step 540 where, based on information contained within the CBR request, the resource manager computes the next node and corresponding egress link. The resource table for the egress link is then examined to determine whether there is a "free" resource entry. If the resource entry is not "Free", the process proceeds to step 542, where the node rejects the CBR request and, from there, to end in step 510. If the resource table entry indicates that the unit's resources are "Free" the process proceeds from step 540

to step 544 where the CBR request is forwarded to the next node and the resource manager awaits a response. When a response is received, the process proceeds to step 546 where, if the response from a downstream node is negative, the process proceeds to step 548 where the resource manager passes this rejection back to the requesting party. From step 548, the process proceeds to end in step 510. If the response from the downstream nodes were affirmative, the resource manager provisions the ingress and egress tributaries as pass-through, returns an acceptance to the requesting party and updates the ingress and egress resource tables to indicate that the status is "Pass", and the Free Flag is set equal to "I". From step 550, the process proceeds to end in step 510 (specification, pages 17-18).

The flowchart of Figure 5 illustrates in greater detail the process by which the resource manager 202, through connection admission control, determines the manner in which ATM traffic is to be switched, should the incoming service request indicate that the traffic is rt-VBR traffic. In one aspect of the invention, the resource manager may be instructed to route all rt VBR requests to either an STM switch or an IP switch, at the option of a service provider, for example. Otherwise, the resource manager 202 may determine whether rt-VBR traffic should be served through a circuit path (that is, through an STM switch fabric within the hybrid switch). This determination may be based on an internal system parameter  $\phi$ , a packet rate ratio threshold, as follows:

employ an STM switch for rt-VBR traffic if  $PPR/SPR \leq \phi$

where:

PPR = peak packet rate

SPR = sustained packet rate

$\phi$  = packet rate ratio threshold

If the rt-VBR traffic is served by an IP switch, the effective bandwidth required for the request may be computed on the request's associated service

contract. The service contract will typically include quality of service requirements, such as delay, packet loss ratio, and the nature of the traffic, which might include, average rate, peak rate, and burst period (specification, pages 18-19).

Requests for rt-VBR traffic may be handled as described in relation to the flow chart of Figure 6, in which the process begins in step 600. In step 600 it is assumed that the traffic request has already been determined to be an rt-VBR request. The process proceeds to step 601, where the resource manager determines whether a service provider has elected to switch all rt-VBR traffic in the same manner as CBR traffic. This election may be evidenced by a flag that has been set, for example. If the service provider has elected to treat all rt-VBR traffic as CBR traffic, that is, in the manner described in the discussion related to Figure 5, the process proceeds to step 604, where the procedures set forth in the discussion related to Figure 5 are employed to switch the traffic. If the service provider has not elected in this manner to treat all rt-VBR traffic as CBR traffic, the process proceeds to step 602, where the resource manager determines whether the rt-VBR traffic is to be switched in a conventional ATM switching manner, such as a linear approximation ATM switching process, or it is to be switched in the same manner as CBR traffic. The resource manager may make this determination using the peak-to-sustained packet rate ratio, opting to switch vt-VBR traffic in the same manner as CBR traffic if the ratio is less than a predetermined threshold value, for example.

If the rt-VBR traffic is to be switched as CBR traffic, that is, in the manner described in the discussion related to Figure 5, the process proceeds to step 604, where the procedures set forth in the discussion related to Figure 5 are employed to switch the traffic. After switching the traffic in this manner, the process proceeds to end in step 606. If, in step 602, the resource manager 202 determines that vt-VBR traffic is to be switched using conventional ATM switching techniques, the process proceeds to step 608, where the switching takes place. From step 608, the process proceeds to end in step 606. IP traffic that has been mapped into other ATM

service categories may be switched using conventional ATM switching techniques (specification, pages 19-20).

Appellant respectfully notes that the above summary of the invention, including any indication of reference numerals, drawings, figures, paragraphs, page numbers, etc. (collectively referred to as "descriptions" of the application) have been provided solely to comply with the U.S. Patent and Trademark Office's rules concerning the appeal of the claims of the present application. As such, the descriptions above are merely exemplary and should not be construed to limit the claims of the present application in any way whatsoever.

**VI. GROUND OF REJECTION TO BE REVIEWED ON APPEAL:**

Appellant seeks the Board's review of the rejection of claims 1 and 12 under 35 U.S.C. §102(e) and claims 1-10, 12-23, 25-31, 33-38 and 40-42 under 35 U.S.C. §103(a).

**VII. ARGUMENTS:**

**The Section 102 and 103 Rejections**

Claims 1 and 12 were rejected under 35 U.S.C. §102(e) as being anticipated by Brueckheimer, U.S. Patent No. 6,574,224 ("Brueckheimer").

Claims 2-6, 10, 13-17, 23, 31 and 38 were rejected under 35 U.S.C. §103(a) as being unpatentable over Brueckheimer in view of Castellano [sic], U.S. Patent No. 6,674,750.

Claims 7 and 18 were rejected under 35 U.S.C. §103(a) as being unpatentable over Brueckheimer and Castellano and in further view of Ko, U.S. Patent No. 5,479,407 ("Ko").

Claims 8 and 19 were rejected under 35 U.S.C. §103(a) as being unpatentable over Brueckheimer, Castellano, Ko and in further view of Caldara, U.S. Patent No. 5,982,771 ("Caldara").

Claims 9 and 25-27 were rejected under 35 U.S.C. §103(a) as being unpatentable over Brueckheimer in view of Castellano, Ko, Caldara and in further view of Houji, U.S. Patent No. 5,832,197 ("Houji").

Claims 20-22 were rejected under 35 U.S.C. §103(a) as being unpatentable over Brueckheimer, in further view of Castellano and Houji.

Claims 28 and 33 were rejected under 35 U.S.C. §103(a) as being unpatentable over Brueckheimer, Castellano and in further view of Ko.

Claims 29 and 34 were rejected under 35 U.S.C. §103(a) as being unpatentable over Brueckheimer in view of Castellano and Ko and Caldara.

Claims 30 and 40-42 were rejected under 35 U.S.C. §103(a) as being unpatentable over Brueckheimer in view of Castellano, Ko, Caldara and Houji.

Finally, claims 35-37 were rejected under 35 U.S.C. §103(a) as being unpatentable over Brueckheimer, Castellano, Ko and in further view of Houji.

Because each of the Section 102 and 103 rejections relies on the disclosure, in full or in part, of Brueckheimer, it is to this reference that Appellant addresses his comments.

Each of the claims of the present invention requires, among other things, the routing of IP traffic to a circuit switch fabric or packet switch fabric depending on an ATM service category of the IP traffic. It is respectfully submitted that Brueckheimer does not disclose or suggest such routing. Though the Office Action indicates that FIG. 1, along with column 5, line 51 to column 6, line 32, 46 and 67 and column 7, lines 1-25 disclose such routing, Appellant respectfully disagrees.

The Appellant has read these excerpts from Brueckheimer. Contrary to the statements made in the Office Action, Appellant was not able to find a disclosure or even a suggestion that Brueckheimer is routing traffic to either a circuit switch or packet switch fabric depending on an ATM service category of IP traffic. Instead, Brueckheimer appears to disclose that IP traffic (e.g., AAL2) is always routed to an ATM switch. That is to say, there is no disclosure or

suggestion in Brueckheimer of a controller that is capable of examining an ATM service category and then routing traffic to either a circuit switch fabric or a packet switch fabric based on the examined service category.

In the Advisory Action, the Examiner reiterates his earlier positions and again recites the same excerpts from Brueckheimer.

For example, the Examiner cites to FIG. 1 of Brueckheimer as depicting the routing of traffic to either a voice/circuit switch or data/packet switch. While this may be accurate, there is no disclosure in Brueckheimer that such routing involves an ATM service category. Rather, voice traffic, at any service category, is routed to a voice switch and data traffic, at any service category, is routed to a data switch.

The Examiner then cites to excerpts from columns 5, 6 and 7 of Brueckheimer to support switching based on an ATM service category. Again, as pointed out above, the switching that is disclosed in Brueckheimer is not based on a service level. Instead it is based on content; voice traffic is routed to a voice switch while data traffic is routed to a data switch.

Because each of the §102 and §103 rejections relies on Brueckheimer, it is respectfully submitted that each of the rejections are inappropriate because Brueckheimer does not disclose or suggest the routing of IP traffic to either a circuit switch fabric or a packet switch fabric depending on an ATM service category of the IP traffic nor do any of the additional references overcome this deficiency.

Accordingly, Appellant respectfully requests that the members of the Board reverse the decision of the Examiner and allow claims 1-42.

#### **VIII. CONCLUSION:**

Appellant respectfully requests the Board to reverse the Examiner's rejection of claims 1-42.

APPELLANT'S BRIEF ON APPEAL UNDER 37 C.F.R. §41.37

U.S. Application No.: 09/620,052


Atty. Docket: 29250-000971/US

The Commissioner is authorized in this, concurrent, and future replies, to charge payment or credit any overpayment to Deposit Account No. 08-0750 for any additional fees required under 37 C.F.R. § 1.16 or under 37 C.F.R. § 1.17; particularly, extension of time fees.

Respectfully submitted,

HARNESS, DICKEY, & PIERCE, P.L.C.

By:



\_\_\_\_\_  
John E. Curtin, Reg. No. 37.602  
P.O. Box 8910  
Reston, Virginia 20195  
(703) 668-8000

JEC:psy

**CLAIMS APPENDIX**

**Claims 1-42 on Appeal:**

1. (Previously Presented) A hybrid telecommunications switch comprising:

at least one circuit switch fabric;

at least one packet switch fabric; and

a controller configured to route IP traffic to the circuit switch fabric or packet switch fabric, depending on an ATM service category of the IP traffic.

2. (Previously Presented) The switch of claim 1 wherein a portion of the circuit switch fabric resources are provisioned for circuit switched traffic and the remaining portion of the circuit switch fabric resources are allocated to IP traffic as the controller routes IP traffic to the circuit switch fabric.

3. (Previously Presented) The switch of claim 2 wherein the controller is further configured to allocate circuit switch fabric resources to traffic falling within an ATM service category.

4. (Previously Presented) The switch of claim 3 wherein the controller is further configured to route IP traffic associated with a constant bit rate (CBR) ATM service category to the circuit switch fabric.

5. (Previously Presented) The switch of claim 3 wherein the controller is further configured to route IP traffic associated with a real time variable bit rate (rt-VBR) ATM service category to the circuit switch fabric.

6. (Previously Presented) The switch of claim 3 wherein the controller is further configured to route IP traffic associated with an ATM traffic service category which is neither CBR nor rt-VBR traffic to the IP switch fabric.

7. (Previously Presented) The switch of claim 3 wherein the controller is further configured to allocate available circuit switch resources, as indicated by a resource table, to received IP traffic requests.

8. (Previously Presented) The switch of claim 7 wherein the controller is further configured to maintain circuit switch ingress and egress resource tables.

9. (Previously Presented) The switch of claim 7 wherein the controller is further configured to pass an IP traffic request to a destination node and to establish an IP traffic path after having determined that all nodes along the proposed traffic path have accepted an IP traffic request.

10. (Previously Presented) The switch of claim 9 wherein the controller is further configured to determine whether IP traffic associated with the rt-VBR ATM service category is to be routed through the circuit switch fabric.

11. (Previously Presented) The switch of claim 10 wherein the controller is further configured to determine whether IP traffic associated with the rt-VBR ATM service category is to be routed through the circuit switch fabric based upon a peak to sustained packet rate ratio.

12. (Previously Presented) A method for routing telecommunications traffic in a hybrid telecommunications switch comprising at least one packet switch fabric, at least one circuit switch fabric, and a controller, including the step of:

routing IP traffic to the circuit switch fabric or packet switch fabric, depending on an ATM service category of the IP traffic.

13. (Previously Presented) The method of claim 12 further comprising the step of:

provisioning a portion of the circuit switch fabric resources for circuit switched traffic, and

allocating the remaining portion of the circuit switch fabric resources to IP traffic as a controller routes the IP traffic to the circuit switch fabric.

14. (Previously Presented) The method of claim 13 further comprising the step of:

allocating circuit switch fabric resources to IP traffic.

15. (Previously Presented) The method of claim 13 further comprising the step of:

routing IP traffic associated with a CBR ATM service category to the circuit switch fabric.

16. (Previously Presented) The method of claim 13 further comprising the step of:

routing IP traffic associated with an rt-VBR ATM service category to the circuit switch fabric.

17. (Previously Presented) The method of claim 13 further comprising the step of:

routing IP traffic not associated with the CBR or rt-VBR ATM service categories to the IP switch fabric.

18. (Previously Presented) The method of claim 13 further comprising the step of:

allocating available circuit switch fabric resources, as indicated by a resource table, to IP traffic requests.

19. (Previously Presented) The method of claim 13 further comprising the step of:

maintaining circuit switch ingress and egress resource tables.

20. (Previously Presented) The method of claim 13 further comprising the step of:

passing an IP traffic request to a destination node.

21. (Previously Presented) The method of claim 20 further comprising the step of:

determining that all nodes along the proposed IP traffic path have allocated circuit switch fabric resources for the IP traffic.

22. (Previously Presented) The method of claim 21 further comprising the step of:

establishing an IP traffic path after the determination step.

23. (Previously Presented) The method of claim 21 further comprising the step of:

determining whether IP traffic associated with the rt-VBR ATM service category is to be routed through an IP switch fabric or a circuit switch fabric.

24. (Previously Presented) The method of claim 23 further comprising the step of:

comparing a sustained packet ratio to a threshold value.

25. (Previously Presented) The switch of claim 9 wherein the controller is further configured to pass an IP traffic request to a destination node.

26. (Previously Presented) The switch of claim 9 wherein the controller is further configured to determine that all nodes along a proposed IP traffic path allocate circuit switch fabric resources for IP traffic.

27. (Previously Presented) The switch of claim 26 wherein the controller is further configured to establish an IP traffic path after determining that all nodes along a proposed IP traffic path allocate circuit switch fabric resources for IP traffic.

28. (Previously Presented) A hybrid telecommunications switch comprising:

at least one circuit switch fabric;

at least one packet switch fabric;

a controller configured to:

route IP traffic to the circuit switch fabric or packet switch fabric, depending on an ATM service category of the IP traffic;

allocate circuit switch fabric resources to traffic falling within an ATM service category; and

allocate available circuit switch resources, as indicated by a resource table, to received IP traffic requests.

29. (Previously Presented) The switch as in claim 28 wherein the controller is further configured to maintain circuit switch ingress and egress resource tables.

30. (Previously Presented) The switch as in claim 28 wherein the controller is further configured to pass an IP traffic request to a destination node and to establish an IP traffic path after having determined that all nodes along the proposed traffic path have accepted an IP traffic request.

31. (Previously Presented) The switch as in claim 30 wherein the controller is further configured to determine whether IP traffic associated with an rt-VBR ATM service category is to be routed through the circuit switch fabric.

32. (Previously Presented) The switch as in claim 31 wherein the controller is further configured to determine whether IP traffic associated with the rt-VBR ATM service category is to be routed through the circuit switch fabric based upon a peak to sustained packet rate ratio.

33. (Previously Presented) A method for routing telecommunications traffic in a hybrid telecommunications switch comprising at least one packet switch fabric, at least one circuit switch fabric, and a controller, including the step of:

routing IP traffic to the circuit switch fabric or packet switch fabric, depending on an ATM service category of the IP traffic;

provisioning a portion of the circuit switch fabric resources for circuit switched traffic;

allocating the remaining portion of the circuit switch fabric resources to IP traffic as a controller routes the IP traffic to the circuit switch fabric; and

allocating available circuit switch fabric resources, as indicated by a resource table, to IP traffic requests.

34. (Previously Presented) The method as in claim 33 further comprising the step of maintaining circuit switch ingress and egress resource tables.

35. (Previously Presented) The method as in claim 33 further comprising the step of passing an IP traffic request to a destination node.

36. (Previously Presented) The method as in claim 33 further comprising the step of determining that all nodes along the proposed IP traffic path have allocated circuit switch fabric resources for the IP traffic.

37. (Previously Presented) The method as in claim 36 further comprising the step of establishing an IP traffic path after the determination step.

38. (Previously Presented) The method as in claim 36 further comprising the step of determining whether IP traffic associated with an rt-VBR ATM service category is to be routed through an IP switch fabric or a circuit switch fabric.

39. (Previously Presented) The method as in claim 38 further comprising the step of comparing a sustained packet ratio to a threshold value.

40. (Previously Presented) The switch as in claim 30 wherein the controller is further configured to pass an IP traffic request to a destination node.

41. (Previously Presented) The switch as in claim 30 wherein the controller is further configured to determine that all nodes along a proposed IP traffic path allocate circuit switch fabric resources for IP traffic.

42. (Previously Presented) The switch as in claim 41 wherein the controller is further configured to establish an IP traffic path after determining that all nodes along a proposed IP traffic path allocate circuit switch fabric resources for IP traffic.





FIGURE 1



200

FIGURE 2

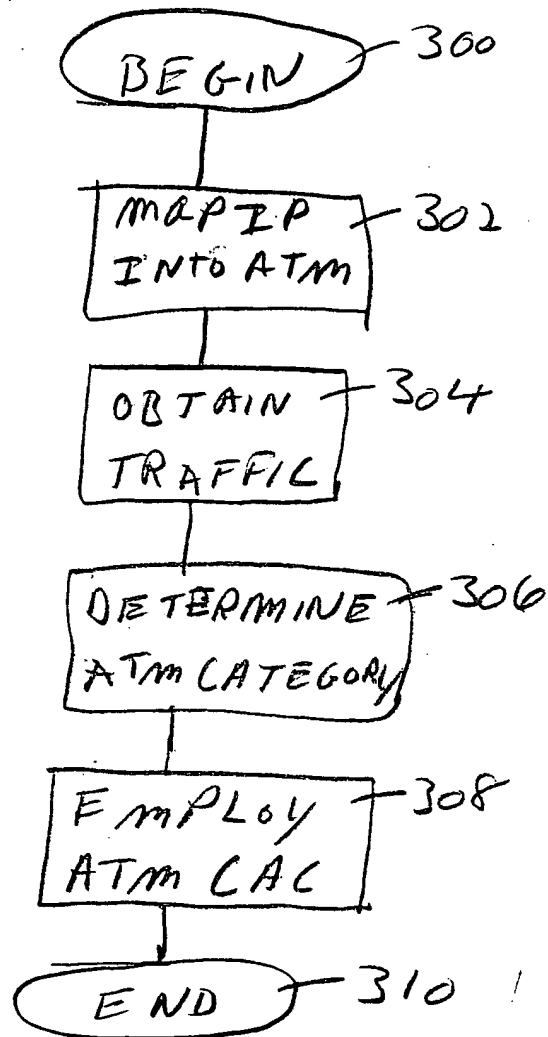


FIGURE 3





FIGURE 4



FIGURES




FIGURE 6